

4

THE FILE COPY

CHEMICAL
RESEARCH,
DEVELOPMENT &
ENGINEERING
CENTER

CRDEC-CR-030

AD-A211 395

IMPROVED PATTERN RECOGNITION TECHNIQUES FOR REAL TIME OPERATION

Gary W. Small


UNIVERSITY OF IOWA
Iowa City, IA 52242

May 1989

DTIC
ELECTE
AUG 14 1989
S D D

This document has been approved
for public release and sale in
distribution is unlimited.

U.S. ARMY
ARMAMENT
MUNITIONS
CHEMICAL COMMAND



Aberdeen Proving Ground, Maryland 21010-5423

89 8 14 055

Disclaimer

The findings in this report are not to be construed as an official Department of the Army position unless so designated by other authorizing documents.

Distribution Statement

Approved for public release; distribution is unlimited.

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE

REPORT DOCUMENTATION PAGE				Form Approved OMB No. 0704-0188	
1a. REPORT SECURITY CLASSIFICATION UNCLASSIFIED			1b. RESTRICTIVE MARKINGS		
2a. SECURITY CLASSIFICATION AUTHORITY			3. DISTRIBUTION/AVAILABILITY OF REPORT Approved for public release; distribution is unlimited.		
2b. DECLASSIFICATION/DOWNGRADING SCHEDULE					
4. PERFORMING ORGANIZATION REPORT NUMBER(S) CRDEC-CR-030			5. MONITORING ORGANIZATION REPORT NUMBER(S)		
6a. NAME OF PERFORMING ORGANIZATION University of Iowa		6b. OFFICE SYMBOL (If applicable)	7a. NAME OF MONITORING ORGANIZATION		
6c. ADDRESS (City, State, and ZIP Code) Iowa City, IA 52242			7b. ADDRESS (City, State, and ZIP Code)		
8a. NAME OF FUNDING/SPONSORING ORGANIZATION CRDEC		8b. OFFICE SYMBOL (If applicable) SMCCR-RSL-S	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER DAAA15-86-C-0034		
8c. ADDRESS (City, State, and ZIP Code) Aberdeen Proving Ground, MD 21010-5423			10. SOURCE OF FUNDING NUMBERS		
			PROGRAM ELEMENT NO.	PROJECT NO.	TASK NO.
11. TITLE (Include Security Classification) Improved Pattern Recognition Techniques for Real Time Operation					
12. PERSONAL AUTHOR(S) Small, Gary W.					
13a. TYPE OF REPORT Contractor		13b. TIME COVERED FROM 86 Apr to 88 Aug		14. DATE OF REPORT (Year, Month, Day) 1989 May	
15. PAGE COUNT 49					
16. SUPPLEMENTARY NOTATION COR: Robert Kroutil, SMCCR-RSL-S, (301) 671-3021					
17. COSATI CODES			18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)		
FIELD	GROUP	SUB-GROUP	Fourier spectroscopy Automatic spectral interpretation Standoff chemical vapor detection algorithm. (file)		
15	03				
17	05				
19. ABSTRACT (Continue on reverse if necessary and identify by block number) Passive Fourier transform infrared remote sensors have shown tremendous potential for use in the automated detection of atmospheric pollutants. To be maximally effective as an early warning alarm device, however, passive infrared sensors must be able to operate in a variety of infrared backgrounds and under conditions in which the background is rapidly changing. For economic reasons, the sensor should be the simplest device that can meet the aforementioned operational requirements. In the work described here, new signal processing strategies have been developed that allow target chemical species to be detected directly from time-domain infrared data. This methodology eliminates the Fourier transform computation through the development of digital filters keyed to the characteristic infrared frequencies of the target molecules. Through this approach, both the computational and instrumental requirements for the sensor are significantly reduced. The performance of this methodology will be illustrated with a variety of test data. <i>Keywords: Dimethyl Methyl Phosphonate, Sulfur Hexafluoride</i>					
20. DISTRIBUTION/AVAILABILITY OF ABSTRACT <input checked="" type="checkbox"/> UNCLASSIFIED/UNLIMITED <input type="checkbox"/> SAME AS RPT. <input type="checkbox"/> DTIC USERS			21. ABSTRACT SECURITY CLASSIFICATION UNCLASSIFIED		
22a. NAME OF RESPONSIBLE INDIVIDUAL SANDRA J. JOHNSON			22b. TELEPHONE (Include Area Code) (301) 671-2914		22c. OFFICE SYMBOL SMCCR-SPS-T

UNCLASSIFIED

UNCLASSIFIED

PREFACE

The work described in this report was authorized under Contract No. DAAA15-86-C-0034. This work was started in April 1986 and completed in August 1988.

The use of trade names or manufacturers' names in this report does not constitute an official endorsement of any commercial products. This report may not be cited for purposes of advertisement.

Reproduction of this document in whole or in part is prohibited except with permission of the Commander, U.S. Army Chemical Research, Development and Engineering Center, ATTN: SMCCR-SPS-T, Aberdeen Proving Ground, Maryland 21010-5423. However, the Defense Technical Information Center and the National Technical Information Service are authorized to reproduce the document for U.S. Government purposes.

This document has been approved for release to the public.

ACKNOWLEDGMENTS

The work described in this report represents the combined efforts of a group of analytical chemists working in the Department of Chemistry at the University of Iowa. The contributions made by Abigail Barber, Joanne Bjerga, Scott Carpenter, Amy Harms, and Thomas Kaltenbach to the work are gratefully acknowledged. Dr. Robert T. Kroutil and coworkers of CRDEC are also acknowledged for supplying the passive FTIR data used in the work and for their participation in many invaluable technical discussions.

Accession For	
NTIS GRA&I	<input checked="checked" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A-1	



Blank

CONTENTS

INTRODUCTION	9
EXPERIMENTAL SECTION	9
RESULTS AND DISCUSSION	10
Overview of Passive FTIR Methodology and Conventional Data Processing	10
Motivation for Alternative Algorithm Development Work	12
Development of Digital Filtering Methodology	18
Mathematical Background	18
Regression Approach to FIR Filter Generation	19
Generation of FIR Filters by Stepwise Regression	20
FIR Matrix Filters	23
Quantitative Comparison of SF ₆ Filter Performance	32
Quantitative Nature of Filter Responses	34
Pattern Recognition Methodology Based on Filtered Interferogram Segments	35
Overview of Methodology	35
Construction of Linear Discriminants	38
Testing of Computed Discriminant	39
Critical Evaluation of Time-Domain Methodology	43
Questions of Spectral Resolution	43
Effects of Phase Errors	43
Source of False Alarms	44
CONCLUSIONS	46
RECOMMENDATIONS	47
REFERENCES	49

LIST OF FIGURES

Figure No.	Page
1.	A 1024-point interferogram collected by a passive FTIR spectrometer mounted on a helicopter 11
2.	The spectrum arising from the application of the FFT to the interferogram depicted in Figure 1 11
3.	Difference spectrum resulting from the subtraction of a background reference spectrum from the sample spectrum of Figure 2 13
4.	A typical laboratory difference spectrum showing the absorption bands of butyric acid superimposed on a flat baseline 13
5.	Segments from six passive FTIR interferograms 15
6.	The region from 900-1000 cm^{-1} is plotted in the transformed spectra corresponding to the interferograms from which the segments in Figure 5 were drawn 15
7.	The frequency-dependent action of a digital filter 16
8.	Three spectra are plotted from 650-1250 cm^{-1} 16
9.	Points 1-50 in the interferograms corresponding to the three spectra in Figure 8 17
10.	Points 50-150 in the interferograms corresponding to the three spectra in Figure 8 17
11.	Points 150-250 in the interferograms corresponding to the three spectra in Figure 8 18
12.	Points 175-250 in the six interferogram segments of Figure 5 after applying the 20-term filter computed by stepwise regression 21
13.	Sum of squares after applying the 20-term filter vs. interferogram number for each interferogram in a typical helicopter data run 21
14.	Sum of squares after applying the 20-term filter vs. interferogram number for each interferogram in a typical ground vehicle data run . 22
15.	Sum of squares vs. interferogram number for the ideal response of a Gaussian filter centered at 941.1 cm^{-1} and having a FWHM of 54.0 cm^{-1} 23
16.	Plot of percentage of variance explained vs. interferogram point ... 25
17.	Plot of number of filter coefficients required vs. interferogram point 25
18.	Plot of percentage variance explained in the matrix filter generation vs. intensity of filtered interferogram points 26

19.	Plot of points 175-250 in six interferogram segments after application of the matrix filter	26
20.	Sum of squares over points 175-250 after applying the matrix filter vs. interferogram number for the helicopter-based data run ..	28
21.	Sum of squares over points 175-250 after applying the matrix filter vs. interferogram number for the ground-based data run	28
22.	Sum of squares over 51 points exceeding the 90% variance cutoff vs. interferogram number for the helicopter-based data run	29
23.	Sum of squares over 51 points exceeding the 90% variance cutoff vs. interferogram number for the ground-based data run	29
24.	Frequency response function for the DMMP matrix filter	30
25.	Sum of squares over points 175-250 vs. interferogram number for the ideal response of the DMMP filter	30
26.	Plot of percentage of variance in the dependent variable explained vs. interferogram point for the DMMP matrix filter	31
27.	Plot of number of filter terms vs. interferogram point for the DMMP matrix filter	31
28.	Sum of squares over 47 points exceeding the 80% variance cutoff vs. interferogram number for the DMMP data run	32
29.	Points 175-250 in six interferograms derived from the helicopter-based data run at 1000 ft	36
30.	The region from 920-970 cm^{-1} in the spectra derived from the interferograms in Figure 29	36
31.	Points 175-250 in the six interferograms after application of the SF_6 matrix filter	37
32.	Plot of the magnitude of the filtered interferogram segment vs. integrated area of the SF_6 spectral band	37
33.	Two-point mean discriminant score vs. interferogram number for a ground-based data run	40
34.	Two-point mean discriminant score vs. interferogram number for a helicopter-based data run	40
35.	Two-point mean discriminant score vs. interferogram number for a helicopter-based data run	41
36.	Two-point mean discriminant score vs. interferogram number for a ground-based data run	41
37.	Plot of percentage variance explained in the matrix filter calculation vs. computed discriminant value for interferogram points 175-250	42

38.	Transformed spectra corresponding to a 76-point segment taken from an interferogram composed of a single cosine frequency (941.1 cm^{-1})	44
39.	Two-point mean discriminant score plot for the helicopter-based data run after introduction of catastrophic phase errors	45
40.	Two-point mean discriminant score plot for the ground-based data run after introduction of catastrophic phase errors	45
41.	Transformed spectrum of interferogram 503 in Figure 35	46

LIST OF TABLES

Table No.		Page
I	Signal-to-Noise Comparison of Filter Performance SF_6 Helicopter-Based Data	33
II	Signal-to-Noise Comparison of Filter Performance SF_6 Ground-Based Data	33
III	Comparison of Computational Efficiency	34

IMPROVED PATTERN RECOGNITION TECHNIQUES FOR REAL TIME OPERATION

INTRODUCTION

Passive infrared spectrometers have come into recent use for remote sensing of atmospheric pollutants¹⁻¹⁰. These devices consist of a Fourier transform infrared (FTIR) spectrometer in which the infrared source has been removed and a set of specialized input optics added. The use of passive FTIR instruments is based on the premise that most gaseous species exhibit a spectral signature in the infrared region that will be superimposed on the normal background infrared emission in the field-of-view of the spectrometer. When controlled by a dedicated computer and appropriate software, the instrument can serve as an automated alarm system for the detection of specific target gases.

Applications for passive FTIR instruments have included both stationary and mobile environments. Recently, increased attention has been given to the mobile application for the following reasons: (1) the instrument is allowed to operate in the cleaner environment of a ground or airborne vehicle; (2) a single instrument can be used to monitor a larger area; and (3) with fewer instruments needed, more highly trained personnel can be used to oversee the installation, maintenance, and operation of the system.

In a mobile application, the spectrometer must be able to operate against a variety of infrared spectral backgrounds and must be able to resist the effects of rapid background changes. In a high-speed vehicle (e.g. helicopter or other aircraft), the data processing tasks must be accomplished rapidly. This includes the complete decision-making procedure for the determination of the presence or absence of the target analyte(s). Additionally, for the system to be practical, the infrared instrumentation must be both reliable and economical. Both reliability and cost are inversely related to the degree of sophistication required in both the instrument and the controlling computer. Further, the required sophistication of both the infrared instrumentation and the computational hardware is directly related to the amount of data that must be collected and the requirements for processing that data.

This report describes efforts at the University of Iowa to reduce both instrumental and computational requirements for passive infrared sensors through the development of novel data processing algorithms. In addition, algorithms are being sought that are resistant to changes in the infrared background. These algorithms are described in detail and evaluated with a variety of test data.

EXPERIMENTAL SECTION

A passive infrared interferometer was constructed by Honeywell Corporation to specifications provided by the U.S. Army Chemical Research, Development, and Engineering Center, Edgewood, MD. The system consisted of a flex-pivot "porch swing" Michelson interferometer based on the design of Walker and Rex¹¹. The spectrometer used a liquid-nitrogen cooled Hg:Cd:Te detector to collect spectral background radiation from the 8 to 12 micron region. For the work reported here, 1024-point interferograms were collected, with a corresponding spectral resolution of approximately 4 cm^{-1} . The standard data collection rate for this instrument was 5 interferograms/sec.

In order to test the feasibility of a mobile application for the sensor and to collect data under a variety of background conditions, the spectrometer was mounted on a shock-absorbing platform and placed in both a helicopter and

a ground vehicle. Several runs of data were collected at a variety of speeds and distances from ground sources of SF₆ and dimethyl methylphosphonate (DMMP). These gases were used to simulate the presence of target analytes that should be detected by the spectrometer. SF₆ has a single strong absorption at 940 cm⁻¹, while DMMP has characteristic bands at 820, 910, and 1040 cm⁻¹. The data collection was performed at Edgewood, MD and Dugway, UT. The collected interferograms were written onto magnetic tape and shipped to the University of Iowa where the algorithm development work was performed.

All computer software used in this work was written in FORTRAN 77 and implemented on a PRIME 9955 interactive computer system operating at the Gerard P. Weeg Computing Center at the University of Iowa. Some computations made use of subroutines from the IMSL statistical package (IMSL, Inc., Houston, TX). Some graphics output made use of the TELAGRAF interactive graphics package (Integrated Software Systems Corp., San Diego, CA). A Hewlett-Packard 7475A digital plotter was used as the output device.

RESULTS AND DISCUSSION

Overview of Passive FTIR Methodology and Conventional Data Processing

In a passive FTIR spectrometer, the initial infrared data is collected in the form of an interferogram, a composite waveform composed of periodic signals corresponding to each of the detected infrared frequencies. In conventional spectral processing, a Fast Fourier Transform (FFT) is applied to the interferogram to decompose it, thereby obtaining the amplitudes of the individual frequency components present. Taken together, these amplitudes define the spectrum of infrared light being detected by the spectrometer. The characteristic infrared frequencies of any chemical compounds present are superimposed on this overall spectrum as either absorption or emission bands. Figure 1 depicts a typical 1024-point FTIR interferogram, and Figure 2 displays the corresponding spectrum yielded by the FFT. The arrow in Figure 2 points to the characteristic band of SF₆, arising here as an absorption (i.e. a decrease in light intensity).

In conventional processing algorithms, the detection of a target gas begins with the collection of a representative background emission interferogram and an interferogram collected with the potential toxic gas source in the field-of-view of the instrument. The FFT is then applied to the collected interferograms. The difference spectrum resulting from subtracting the sample and background spectra is analyzed for the spectral signature of the compound being monitored. Under conditions in which the background is changing, however, the difference spectrum can contain anomalous features, possibly resulting in either false alarms or missed alarms.

As an illustration, the interferogram depicted in Figure 1 was collected when the spectrometer was mounted in a helicopter flying at 25 knots and at an altitude of approximately 200 feet. A reference spectrum was obtained by transforming an interferogram collected two seconds earlier. Based on the position of the helicopter and an inspection of the spectrum, no SF₆ is indicated in this reference spectrum. The spectrum in Figure 2 (i.e. the sample spectrum in this example) and the reference spectrum were each normalized to unit area to correct for overall differences in infrared intensity. Figure 3 depicts the difference spectrum obtained by subtracting the reference spectrum from that of the sample. The SF₆ absorption remains evident, but a number of spectral artifacts have been introduced that may interfere with numerical decision-making algorithms designed to produce a yes or no regarding the presence of SF₆.




Figure 1. A 1024-point interferogram collected by a passive FTIR spectrometer mounted on a helicopter. Light intensity detected is plotted vs. interferogram point number. The interferogram has been rotated such that the maximum intensity (center burst) is at point 1. A ground source of SF_6 was in the field-of-view of the instrument when this interferogram was collected.




Figure 2. The spectrum arising from the application of the FFT to the interferogram depicted in Figure 1. Light intensity is plotted vs. spectral frequency (cm^{-1}). The absorption band due to the presence of SF_6 is indicated by the arrow.

By way of comparison, the results of a typical laboratory spectral subtraction are depicted in Figure 4. A normalized reference spectrum was subtracted from a sample spectrum of butyric acid to produce the difference spectrum shown in the figure. These spectra were collected on an IBM M98 laboratory FTIR spectrometer under stable background conditions. The characteristic absorption bands of butyric acid are clearly evident against a flat baseline. This example clearly demonstrates the severity of the artifacts introduced in Figure 3 by an unstable spectral background.

A variety of techniques have been investigated to help overcome the presence of these artifacts. Digital filters have been employed to attempt to discriminate against the artifacts. In addition, attempts have been made to characterize common background features in hopes of accounting for the presence of background changes. All such methods are inherently limited, however, as an infinite variety of possible background changes exist.

Extensive inspections of spectra collected under moving background conditions reveal one overriding principle: spectral subtraction is a laboratory procedure based on the assumption that a stable spectral background can be established. The use of spectral subtraction techniques under conditions of changing backgrounds is an invalid application of the technique.

Motivation for Alternative Algorithm Development Work

The spectral subtraction procedure described above is useful in a laboratory environment because it represents a conceptually simple way to extract an absorption or emission band from the infrared background response, provided the background response is known. The procedure is applied to spectra because the FFT represents the conceptually simplest way to remove information regarding spectral frequencies not pertinent to the detection of the analyte.

Often, the conceptually simplest approach to a data analysis problem is not the most computationally efficient. The standard Cooley-Tukey FFT requires over 10,000 multiplications for a single 1024-point interferogram. Additionally, as argued in the discussion above, this conceptually simplest approach based on spectral subtraction is not really appropriate for the present application. The goal for our research has been to develop a data processing algorithm for mobile passive FTIR spectrometers that makes better use of the available computational resources, while, at the same time, overcoming the problems of changing spectral backgrounds.

As motivated above, the goal for such an algorithm is twofold. First, the information pertaining to the characteristic frequencies of a target molecule must be separated from the information pertaining to other frequencies. Second, any absorption or emission band present in the frequency window of interest must be extracted from the spectral baseline. This extracted information can then be used with a standard decision-making (pattern recognition) algorithm to produce a yes/no response regarding the presence of the target analyte.

As information regarding the characteristic frequencies of a target molecule is also present in the raw interferogram, it can be argued that use of the Fourier transform is an inefficient strategy for frequency selection. The problem with an interferogram-based analysis, however, centers on the composite nature of the waveform. The periodic signals corresponding to the individual frequencies co-add to form the interferogram. Information regarding the specific frequencies of interest is therefore obscured by frequency information that has no bearing on determining the presence of the target compound.

This problem is illustrated by an inspection of Figure 5. Segments (points 175-250 in the 1024-point rotated interferogram) from six




Figure 3. Difference spectrum resulting from the subtraction of a background reference spectrum from the sample spectrum of Figure 2. Difference in intensity is plotted vs. frequency (cm^{-1}). The artifacts introduced are due to changes in the background between the times the two spectra were collected.




Figure 4. A typical laboratory difference spectrum showing the absorption bands of butyric acid superimposed on a flat baseline. Difference in intensity is again plotted vs. frequency (cm^{-1}). This type of spectrum results when a stable spectral background exists.

interferograms are plotted. Figure 6 depicts the region of 900-1000 cm^{-1} in the corresponding transformed and normalized spectra. Counting from the top of the figure, spectra 3, 5, and 6 appear to contain SF_6 information, although the bands are very weak. It is apparent from an inspection of Figure 5, however, that it would be difficult, at best, to base a detection procedure on the raw interferogram information alone.

Time-domain digital filters¹² are mathematical transforms that operate on time-domain data in a frequency-dependent manner. The result of the transformation is a filtered time-domain signal that has had information pertaining to certain frequencies suppressed. The frequency dependence of a filter is encoded in its frequency response function. In effect, this function is a spectrum that defines which frequencies the filter will pass. Figure 7 demonstrates the action of such a filter in the spectral domain. The upper plot in the figure is the same passive FTIR spectrum depicted in Figure 2. The middle plot is a filter frequency response function centered on the SF_6 absorption band. This is a Gaussian-shaped function with a full-width at half-maximum (FWHM) of 71.1 cm^{-1} . The result of applying this filter to the FTIR spectrum is depicted at the bottom of the figure. The filter effectively extracts the SF_6 band from the spectral baseline.

The above discussion focused on the familiar spectral domain. The advantage of a time-domain filter, however, is that a filter with the frequency response depicted in Figure 7 can operate directly on the interferogram, producing a filtered interferogram whose transformed spectrum is identical to that shown at the bottom of the figure. In this manner, the Fourier transform has been avoided and the spectral band of interest has been extracted from the baseline.

Figures 8-11 demonstrate graphically that an interferogram-based detection scheme is a feasible, straightforward approach to the passive FTIR detection problem. Figure 8 depicts, from top to bottom, a Gaussian-shaped filter frequency response, the spectrum resulting from applying this filter to a synthetic Lorentzian absorption band, and the spectrum resulting from applying the filter to a synthetic Lorentzian emission band. The two Lorentzian bands are identical in shape, width, and area.

If the inverse FFT is applied to the three spectra, the corresponding interferograms are obtained. Figures 9-11 depict points 1-50, 50-150, and 150-250, respectively, in these three computed interferograms. An inspection of these figures reveals three fundamental facts that allow an interferogram-based detection scheme to work.

First, a separation of information exists in the interferogram based on the width of the spectral feature. The three segments in Figure 9 are identical. This indicates that the bulk of the information pertaining to the Gaussian filter function is contained in the first 50 interferogram points. The time-domain form of the Gaussian has effectively damped out by point 110. The information from points 110-250 pertains almost entirely to the Lorentzian band. This observation implies that only a short segment of the interferogram may be necessary in order to detect the presence of a particular spectral band. By collecting only a short segment (termed windowing the interferogram), instrumental requirements are reduced and many background artifacts are automatically eliminated.

Second, the only difference between the absorption and emission spectra exists in the region of points 50-100. Beyond point 100, the signals are identical in magnitude, differing only in phase. Stated differently, the lower two plots in Figure 11 are exact mirror images. This observation confirms that absorption and emission bands can be detected in the interferogram with equal ease.

Third, information regarding the intensity of the absorption or emission band is encoded in the intensity of the periodic interferogram signal produced




Figure 5. Segments from six passive FTIR interferograms. Light intensity is plotted vs. interferogram point for points 175-250.




Figure 6. The region from 900-1000 cm^{-1} is plotted in the transformed spectra corresponding to the interferograms from which the segments in Figure 5 were drawn. Small SF_6 bands around 940 cm^{-1} are evident in the third, fifth, and sixth spectra, counting from the top of the figure.




Figure 7. The frequency-dependent action of a digital filter. From top to bottom: a typical passive FTIR spectrum, a filter frequency response function, the result of applying the filter to the spectrum.




Figure 8. Three spectra are plotted from 650-1250 cm^{-1} . Top: A Gaussian filter frequency response. Middle: The result of applying this filter to extract a Lorentzian absorption band. Bottom: The result of extracting an identical Lorentzian emission band.




Figure 9. Points 1-50 in the interferograms corresponding to the three spectra in Figure 8.




Figure 10. Points 50-150 in the interferograms corresponding to the three spectra in Figure 8.




Figure 11. Points 150-250 in the interferograms corresponding to the three spectra in Figure 8.

by the band. This suggests that a possible qualitative indicator of the presence of an absorption or emission may be a simple power calculation (e.g. sum of squares of intensity) over an appropriate interferogram segment. Moreover, such a computed value may have quantitative spectral information also.

The above discussion demonstrates unequivocally that the combination of windowing the interferogram and applying a time-domain digital filter can be used to detect the presence of specific absorption or emission bands. Through this procedure, the band information can be extracted from the infrared background without the need for the Fourier transform or for any type of spectral subtraction or background characterization. The development and application of this methodology to passive FTIR data will now be discussed.

Development of Digital Filtering Methodology

Mathematical Background. In the discussion above, the frequency-domain form of a digital filter was applied to a spectrum by multiplying the spectrum by the frequency response function of the filter. This calculation has a corresponding form in the time domain that is expressed by the convolution integral

$$y(t) = \int_{-\infty}^{\infty} h(k)x(t-k) dk = H(f)X(f) \quad (1)$$

where $y(t)$ is the filtered interferogram, the product, $H(f)X(f)$, is the filtered spectrum, $H(f)$ is the frequency response of the filter, $X(f)$ is the transformed spectrum, $x(t)$ is the raw interferogram, and $h(t)$ is the inverse Fourier transform of the frequency response, termed the impulse response of the filter. The terms, $H(f)$ and $X(f)$, are functions of frequency, f , while the terms, $y(t)$, $x(t)$, and $h(t)$, are functions of the time variables, t or k .

The above equation pertains to continuous functions. For digitized data, the discrete convolution integral has the form

$$y_t = \sum_{k=-\infty}^{\infty} h_k x_{t-k} \quad (2)$$

where y_t is one point on the filtered interferogram, and h_k and x_{t-k} define discrete points on the impulse response function and interferogram, respectively. Given that the interferogram is only sampled over a finite range, the discrete form of the convolution integral must be approximated by use of a finite series of terms. This approximation most often takes the form

$$y_t = h_0 x_t + \dots + h_n x_{t-n} \quad (3)$$

where $n+1$ terms are summed to estimate each filtered interferogram point, y_t . The terms, h_k , can be considered weighting coefficients that determine the frequency dependence of the filter. Effectively, each filtered interferogram point is formed from a linear combination of the corresponding raw data point and a set of preceding raw data points. This type of filter, defined by the h_k , is termed a finite impulse response (FIR) filter.

The h_k terms must be found that best approximate the action of the filter frequency response function. This is most often accomplished by forming a polynomial series approximation to the frequency response. By approximating the frequency response with a finite series of terms, the h_k can be computed by direct calculation. The Remez exchange method^{13,14} is perhaps the most commonly used algorithm for generating a finite series approximation of the frequency response function.

As discussed previously, the passive FTIR application requires the development of narrow-bandpass filters that can isolate specific spectral bands. Unfortunately, as the bandpass narrows, it takes increasingly more terms to approximate the frequency response. This causes the corresponding time-domain filter to become increasingly large in terms of the number of h_k required. Filters generated for the passive FTIR application by this conventional approach often required as many as 100 terms. Even if a 100-term filter is only applied to 100 interferogram points, 10,000 multiplications are required. This results in no savings over applying the FFT to the entire interferogram. For this reason, it was judged essential that methodology be developed for reducing the number of required filter terms.

Regression Approach to FIR Filter Generation. As noted, eq. 3 is a linear model that relates a filtered interferogram point, y_t , to a series of raw interferogram points, x_{t-n} . Regression analysis is the branch of statistics that focuses on techniques for the construction and evaluation of such models. In regression terms, the y_t define a dependent variable, while the x_{t-n} define a set of independent variables. The h_k are regression coefficients.

It was judged potentially fruitful to investigate the application of regression techniques to the construction of FIR filters. Two motivating factors led to this conclusion. First, in any regression model, some independent variables are always more statistically significant than others. These significant variables contribute most to the explanation of variance in the dependent variable. It was hypothesized that if regression procedures could be used to estimate the h_k , it may be possible to delete some of the x_{t-n} terms as being insignificant. The resulting filters would thus have fewer terms. Second, an inspection of eq. 1 reveals that, in a conventional FIR filter, the h_k depend only on the frequency response function. In one sense, this is an advantage in that a set of h_k can be used to filter any type of time-domain signal and any segment of that signal. For the present application, however, the only time-domain signal encountered is the interferogram produced by the interferometer and detector of the passive FTIR

spectrometer. Moreover, as demonstrated previously, it should be possible to filter only a short segment of that interferogram. It was hypothesized that fewer filter terms would be required if a regression approach could be used to derive a set of h_k that are dedicated to filtering specific segments of passive FTIR interferograms.

Generation of FIR Filters by Stepwise Regression. The generation of an FIR filter via regression methodology begins as before with the definition of the frequency response of the filter. For example, to generate an SF₆ filter, a Gaussian function would typically be defined, centered at the SF₆ peak maximum and possessing a certain width. The exact position and width of the bandpass define parameters to be optimized. As noted previously, if this Gaussian function is multiplied by a sample spectrum, a filtered spectrum results. The inverse Fourier transform of this filtered spectrum is a corresponding filtered interferogram. If the desired time-domain filter were available, the action of the filter on the sample interferogram would produce a filtered interferogram identical to that obtained through the Fourier transform/Gaussian multiplication/inverse transform step outlined above. Thus, points in this generated filtered interferogram define the dependent variable for the regression analysis. The specific points used would be those defining the interferogram segment for which the dedicated filter is desired. The selection of which points to use reflects an additional parameter to be optimized.

A standard multiple linear regression analysis can thus be performed to derive the desired h_k . In this calculation, a pool of potential independent variables can be used at $k=0, \dots, n_{\max}$. These variables would be derived from points in the same raw interferogram whose transformed spectrum was used in the generation of the dependent variable. The specific interferogram selected for the computation does not appear to matter, although it has been common practice throughout this work to base the calculation on an interferogram whose spectrum contains the spectral band that is serving as the target of the filter development.

A stepwise regression algorithm¹⁵ is used to select those x_{t-k} that contribute most to the model. This algorithm begins by selecting the single independent variable that has the highest correlation with the dependent variable. A one-term model is then formed based on this selected variable. Consecutive terms are added to the model in a stepwise manner. At each step, the variables remaining in the pool are evaluated for their correlation with the variance in the dependent variable that has not been explained by the terms previously selected. The variable chosen through this procedure is added to the model, and the process is repeated until no remaining variables in the pool meet a minimum standard of correlation. This correlation test is typically referenced to a statistical E-distribution.

As an example, an SF₆ filter for interferogram points 175-250 was generated based on a Gaussian frequency response centered at 941.1 cm⁻¹ and possessing a FWHM of 54.0 cm⁻¹. An interferogram collected during a ground-based data run was used in the regression procedure. The interferogram was normalized based on the sum of squares of points 175-250. This normalization method, based on the selected segment, is the best method found for correcting for differences in overall infrared energy between interferograms. For the stepwise regression, a pool of variables was used from $k=0$ to 70. To be selected, variables had to be significant at the 95% level, based on the E-distribution. The regression procedure yielded a 20-term filter. The correlation coefficient for the regression was 0.992, corresponding to 98.39% (100 times 0.992 squared) of the variance in the dependent variable being explained by the derived model.

Figure 12 presents the results of applying this filter to the six raw interferogram segments depicted previously in Figure 5. Interferogram




Figure 12. Points 175-250 in the six interferogram segments of Figure 5 after applying the 20-term filter computed by stepwise regression.




Figure 13. Sum of squares after applying the 20-term filter vs. interferogram number for each interferogram in a typical helicopter data run.




Figure 14. Sum of squares after applying the 20-term filter vs. interferogram number for each interferogram in a typical ground vehicle data run.

segments 3, 5, and 6, thought to contain SF_6 information, are now clearly different from the remaining segments, thought to contain no SF_6 information.

Figures 13 and 14 result from applying this filter to each interferogram in two experimental runs. The interferogram used in the generation of the filter was not taken from either of these two runs. Before applying the filter, each interferogram was normalized based on the sum of squares of points 175-250. After filtering points 175-250 in each interferogram, the sum of squares was computed across the range of filtered points. In the figures, the computed sums of squares are plotted vs. interferogram number. As motivated by the discussion of Figures 9-11, the peaks in Figures 13 and 14 correspond to interferograms in which SF_6 information is present. The interferograms used in the generation of Figure 13 were collected with the spectrometer mounted on a helicopter flying at 80 knots and at an altitude of 1000 feet. The spectrometer was mounted on a ground vehicle moving at 10 mph during collection of the data depicted in Figure 14. Both data sets used represent worse-than-average helicopter and ground-based runs in terms of the strength of the SF_6 signals present. Much higher signal-to-noise is always observed in the helicopter runs, as the spectrometer is looking at the target gas against the earth, a strong infrared radiator.

The center of the filter bandpass, the filter width, and the interferogram segment used are parameters that must be optimized for each spectral band investigated. The parameter values used above were selected based on evaluations of what we term ideal filter responses. Across a data run, this procedure involves Fourier transforming each interferogram, multiplying the computed spectrum by the desired filter frequency response, and then inverse transforming the filtered spectrum back to the interferogram domain. The sum of squares is then computed on this ideally filtered interferogram. A plot of the sums of squares vs. interferogram number can then be evaluated in terms of signal-to-noise. This effectively evaluates the

degree of optimization of the filter frequency response and the selected interferogram segment. The frequency response and interferogram segment are sought that maximize this signal-to-noise ratio.

This procedure assumes, of course, that a time-domain version of the filter can be generated. Figure 15 depicts the ideal filter response across the helicopter data set described above. Interferogram normalization based on the sum of squares of points 175-250 was used as before. The frequency response and interferogram segment used were identical to that employed in the generation of the 20-term filter. It is clearly evident from an inspection of the figure that the computed time-domain filter does not duplicate the performance of its frequency-domain counterpart.

The limitations of the stepwise regression procedure were further illustrated when attempts were made to construct a multiple-bandpass filter for DMMP. Very low correlation coefficients were observed in the regression analysis.

FIR Matrix Filters. An inspection of eq. 1 reveals that each point on the filtered interferogram is derived through the separate evaluation of the convolution integral for that point. As noted previously, a time-domain filter must approximate this convolution integral at each point. In a conventional FIR filter, however, the same filter coefficients and the same surrounding points are used to define each filtered point. To achieve the best approximation of the convolution integral, however, it can be argued that a different filter model should be used at each point. For each point, such a filter would have a different set of filter coefficients and be based on a different set of surrounding points.




Figure 15. Sum of squares vs. interferogram number for the ideal response of a Gaussian filter centered at 941.1 cm^{-1} and having a FWHM of 54.0 cm^{-1} .

By deriving separate filters to be used on each interferogram point, it was hypothesized that the overall accuracy of the filtering process could be increased. To perform the derivation of such a filter, however, a set of interferograms must be employed. For example, a filter to be used on interferogram point 175 must be derived through the use of point 175 across a set of interferograms. It is clear that the use of a set of filters may offer increased accuracy. It is also true that such a procedure may be faster computationally. Since each filter is now performing a simpler task, fewer terms in each filter may be needed.

The derivation of a set of filters is a direct extension of the stepwise regression methodology discussed previously. Separate stepwise regressions are performed, one for each interferogram point treated. The result of this process is a set of filters stored as a matrix of filter coefficients. For this reason, we have termed the resultant filter set a "matrix filter."

This methodology was tested by developing a matrix filter for SF₆ based on the same frequency response and interferogram segment used above. Two hundred interferograms from a ground-based data run were used in this calculation. Interferogram normalization based on points 175-250 was used as before. The pool of independent variables used at each point was defined as $k=0$ to 70. To be judged significant, independent variables had to meet a 97% E-criterion. The 97% criterion was used after discovering that a 95% E allowed more terms into the models than were necessary. No performance drop-off was observed in switching to a 97% cutoff. Subsequent re-investigation of the 20-term standard FIR filter revealed that a 95% E cutoff was necessary in that case to maintain filter performance.

In deriving the matrix filter coefficients, it was discovered that when computations are carried out across a set of interferograms, the presence of random low-frequency noise prevents the regression procedure from working properly. For this reason, each interferogram was preprocessed with a 4-term low-frequency cutoff filter to eliminate low-frequency noise.

Figures 16 and 17 depict the results of the regression procedure. Figure 16 is a plot of the percentage of variance in the dependent variable explained by each filter model vs. interferogram point number, while Figure 17 is a plot of the number of filter terms required vs. interferogram point. With only a few exceptions, the percentage of variance explained is high. Fourteen filter terms are needed on average, with the largest filter requiring 21 terms. The plot in Figure 16 reveals a periodic phenomenon that is presumably related to the filter bandpass. As depicted in Figure 12, those points on the waveform passed by the filter very nearly define a single cosine function. The points near zero on this waveform would be expected to have low signal-to-noise. It seems plausible that the regression computation for these points would be dominated by random (i.e. uncorrelated) noise, thus resulting in a lower percentage of variation explained.

To test this hypothesis, the computed matrix filter was applied to the three SF₆-active interferograms discussed previously in relation to Figures 5 and 12. The percentage of variance explained in the filter generation was plotted vs. the corresponding intensity of each filtered point. This plot is depicted in Figure 18. The general trend in the plot is for those points with low percentages of variance explained to be associated with intensities near zero. Correspondingly, there are no points near the extrema (0.04, -0.04) that exhibit low correlation in the filter generation.

The generated matrix filter was tested by applying it to a number of interferograms not included in the regression computation. As an initial test, the filter was applied to the same six interferograms described previously in the discussion of Figures 5 and 12. Figure 19 is a plot of these interferogram after filtering. The SF₆-active interferograms stand out somewhat more cleanly in this plot than in Figure 12.




Figure 16. Plot of percentage of variance explained vs. interferogram point.




Figure 17. Plot of number of filter coefficients required vs. interferogram point.




Figure 18. Plot of percentage variance explained in the matrix filter generation vs. intensity of filtered interferogram points.




Figure 19. Plot of points 175-250 in six interferogram segments after application of the matrix filter. Counting from the top, the third, fifth, and sixth interferogram segments are thought to contain SF_6 information.

The filter was next applied to the helicopter and ground-based data runs discussed in relation to Figures 13 and 14. The interferograms used in the generation of the filter were not part of either of these data sets. The 4-term low-frequency cutoff filter was applied to each interferogram, and the normalization procedure based on points 175-250 was applied. Following application of the filter, the sum of squares was computed across points 175-250. Figure 20 is a plot of the computed sum of squares vs. interferogram number for the helicopter run, and Figure 21 is the corresponding plot for the ground-based run. Figures 20 and 21 can be compared directly to Figures 13 and 14, respectively. From this comparison, it is clear that the matrix filter produces superior results in terms of signal-to-noise.

An additional advantage of the matrix filter approach is found in the point-by-point feedback from the regression procedure. Figure 16 indicates that some points can be filtered more reliably than others, based on the success of the filter generation procedure. An experiment was conducted in which the 25 points below the 90% level in Figure 16 were eliminated from the sum of squares calculation. Figures 22 and 23 are the resulting sum of squares plots based on the remaining 51 points in the 175-250 segment. Figure 22 depicts the helicopter run, while Figure 23 depicts the ground-based run. The signal-to-noise for these plots appears slightly better than that observed for Figures 20 and 21, validating the point deletion procedure.

The matrix filter approach was also evaluated for its ability to generate multiple bandpass filters. As noted previously, DMMP has three characteristic spectral bands. Based on the ideal filter response procedure outlined for SF₆, the three-band frequency response depicted in Figure 24 was judged optimum. The ideal response for this filter is depicted in Figure 25. Based on an inspection of spectra, interferograms 277-360 are thought to contain DMMP signals. Interferogram points 175-250 were again found to be optimum. One modification to the previous procedure was required, however. In the investigation of the DMMP ideal response, interferogram normalization based on the entire interferogram (points 1-1024) was found to produce markedly better results than normalization based solely on points 175-250. The reason for this discrepancy with the SF₆ results is unclear. Unfortunately, only a very limited amount of passive FTIR DMMP data is available to help pinpoint this phenomenon. This problem is made worse by very low signal-to-noise ratios in the available DMMP interferograms.

The matrix filter for DMMP was computed based on 200 interferograms. The interferograms were normalized based on points 1-1024, and the 4-term low-frequency cutoff filter was applied as before. The pool of independent variables used was defined by $k=0$ to 70. Variables had to meet a 95% F -criterion in order to enter the filter models. The computed filter exhibited significantly lower values for the percentage of variance explained in the dependent variable than the filter computed previously for SF₆. Figure 26 is a plot of percent variance explained vs. interferogram point, while Figure 27 plots the number of filter terms required vs. interferogram point. A more complex structure is seen in Figure 26 than was seen in the corresponding plot for the SF₆ filter, although a definite pattern exists. The filters average nine terms. These models are smaller than those computed for SF₆, as fewer terms in the DMMP calculation met the F -criterion. Relaxing the F -criterion for the entry of variables increases the sizes of the models, but the percentage of variance explained does not increase significantly and the larger models do not exhibit better performance. We feel the low signal-to-noise of the DMMP interferograms is being manifested in the overall lower correlations and the smaller number of statistically significant variables.

Following the procedure found useful for the SF₆ matrix filter, those points exhibiting low percentages in Figure 26 were not included in the sum of squares calculation after application of the filter. An 80% cutoff was chosen




Figure 20. Sum of squares over points 175-250 after applying the matrix filter vs. interferogram number for the helicopter-based data run.




Figure 21. Sum of squares over points 175-250 after applying the matrix filter vs. interferogram number for the ground-based data run.




Figure 22. Sum of squares over 51 points exceeding the 90% variance cutoff vs. interferogram number for the helicopter-based data run.




Figure 23. Sum of squares over 51 points exceeding the 90% variance cutoff vs. interferogram number for the ground-based data run.




Figure 24. Frequency response function for the DMMP matrix filter.




Figure 25. Sum of squares over points 175-250 vs. interferogram number for the ideal response of the DMMP filter.




Figure 26. Plot of percentage of variance in the dependent variable explained vs. interferogram point for the DMMP matrix filter.




Figure 27. Plot of number of filter terms vs. interferogram point for the DMMP matrix filter.




Figure 28. Sum of squares over 47 points exceeding the 80% variance cutoff vs. interferogram number for the DMMP data run.

arbitrarily after an inspection of the figure. Forty-seven points remained after applying this cutoff. Figure 28 is a plot of the sum of squares over these points after filtering vs. interferogram number for a ground-based DMMP data run. Due to the small quantity of DMMP data available, the interferograms used in the filter generation were derived from this data set. The overall signal-to-noise ratio of this plot is greater than that of the ideal plot. We believe this is due to the identification and deletion of low signal-to-noise points from the sum of squares calculation.

Quantitative Comparison of SF₆ Filter Performance. The filtering results presented above are best compared in a quantitative manner. This comparison was made based on two criteria: (1) signal-to-noise ratio of the sum of squares plots and (2) computational efficiency. Based on inspections of transformed spectra, SF₆-active regions were established in both the helicopter and ground-based data runs. A baseline set of 200 interferograms was selected in each data set from among those interferograms judged to contain no SF₆. For each plot, the sum of squares values for these baseline interferograms were used to compute a standard RMS noise level. In addition, simple linear regression was used to define a baseline model for each plot. For each interferogram in the SF₆-active regions, the baseline model was used to compute a predicted baseline value, and the difference was computed between the sum of squares value for the interferogram and the baseline value. The ratio of this difference to the RMS noise was taken as the signal-to-noise ratio. This computation corrects for sloping baselines.

Table I is a signal-to-noise comparison of the plots produced by the standard and matrix FIR filters for the SF₆ helicopter-based data run. The signal-to-noise computation described above was used in the generation of the table entries. Five SF₆-active regions were found in this data set, corresponding to interferograms 231-246, 320-333, 419-427, 532-535, and 599-605. Table II presents the corresponding information for the SF₆ ground-based run. One SF₆-active region was found, corresponding to interferograms 358-387.

Table I
Signal-to-Noise Comparison of Filter Performance
SF₆ Helicopter-Based Data

(Based on Sum of Squares After Filtering)

Filter	Signal-to-Noise Ratio for Interferogram Region				
	231-246	320-333	419-427	532-535	599-605
Ideal response	73.48	79.30	57.39	47.62	68.02
20-term standard FIR	34.06	36.33	26.46	23.68	30.37
Matrix FIR, 76 points	61.41	59.25	52.66	38.85	69.32
Matrix FIR, 51 points based on 90% cutoff	81.68	78.47	69.53	49.39	90.82

Table II
Signal-to-Noise Comparison of Filter Performance
SF₆ Ground-Based Data

(Based on Sum of Squares After Filtering)

Signal-to-Noise for Interferograms	
Filter	358-387
Ideal response	11.34
20-term standard FIR	4.69
Matrix FIR, 76 points	19.55
Matrix FIR, 51 points based on 90% cutoff	22.90

An inspection of Tables I and II serves to reinforce the conclusions reached from a visual inspection of the filtering results. The matrix FIR filter significantly outperforms the standard FIR filter. This is particularly evident in the low signal-to-noise data found in the ground-based data run. Furthermore, the deletion of poorly modeled points from the sum of squares calculation significantly enhances the signal-to-noise ratios of the plots. The matrix filter methodology is thus judged superior in terms of its ability to extract low-signal spectral information.

One unexplained phenomenon is seen in the matrix filter producing higher signal-to-noise than the corresponding ideal response for the ground-based data. This can be explained when points are deleted from the matrix filter sum of squares calculation based on the argument that the deleted points contain little signal information. For the ground-based data, however, the matrix filter response over points 175-250 has higher signal-to-noise than the ideal response over the same points. This phenomenon remains under investigation.

Table III compares the standard and matrix filter approaches to the FFT in terms of the number of multiplications required. Also indicated is the percentage reduction in number of multiplications gained when each of the filter approaches is used in place of the FFT. The matrix filter based on 51 points is the fastest of the three filters. Its use in place of the FFT effectively results in a 90% savings in the number of multiplications required. These results further validate the selection of the matrix filter based on selected interferogram points as the best available filtering methodology for the passive FTIR application.

Quantitative Nature of Filter Responses. The discussion of Figures 8-11 indicated that quantitative information may be present in the intensity of the filtered interferogram segments. Conventional thinking in passive FTIR has

Table III
Comparison of Computational Efficiency

Operation	Number of Multiplications	% Savings Over FFT
1024-point Cooley-Tukey FFT	10240	0.0
20-term standard FIR on points 175-250	1520	85.2
*SF ₆ Matrix FIR on points 175-250	1394	86.4
*SF ₆ Matrix FIR on 51 points based on 90% cutoff	1065	89.6

*Includes multiplications for 4-term low-frequency cutoff filter.

held that quantitative determinations are not possible. For this reason, emphasis has been placed on the development of active infrared sensors based on CO₂ lasers. The ranging capability of the laser systems allows a remote concentration-path length (CL) determination to be made for a target analyte.

The placement of a passive FTIR sensor in an airborne vehicle makes quantitative determinations feasible, however. If the input optics of the spectrometer are directed at the ground, effectively all of the infrared energy entering the interferometer arises from ground emission. For an aircraft in level flight, this translates to a fixed path length for the spectroscopic determination, directly analogous to a fixed-source laboratory experiment. Moreover, the altimeter of the aircraft provides a direct measure of the path length.

It would thus seem feasible to develop quantitative calibrations for target analytes. Toward that end, we have evaluated the quantitative information present in filtered interferogram segments. Figure 29 is a plot of points 175-250 in six interferograms derived from the helicopter-based data run at 1000 ft. Figure 30 depicts the region of 920-970 cm⁻¹ in the transformed spectra arising from these interferograms. No SF₆ can be detected in the topmost spectrum, but increasing levels of SF₆ are present in the lower spectra. Figure 31 is a plot of points 175-250 after application of the SF₆ matrix filter. The intensity of the filtered signals increases from the top to bottom of the plot in agreement with the observed spectra.

To evaluate this apparent quantitative information, the spectral bands in Figure 30 were integrated by estimating the baseline across each band and integrating between the baseline and the band. The magnitude of each filtered segment was computed by taking the square root of the sum of the squares of the intensities. Figure 32 is a plot of the computed magnitudes of the filtered segments vs. the computed areas of the corresponding spectral bands. The individual points are shown, as well as the regression line derived from the points. The plot is highly linear. Linear regression reveals that 99.0% of the variance in the dependent variable is explained by the derived linear model. The computed *t*-value for the significance of the slope of the plot is 19.47. A *t*-value greater than 4.0 is considered strong. In addition, significant dynamic range is available in the plot between the smallest SF₆ band and the baseline point (no SF₆). These results confirm that a quantitative calibration can be obtained directly from information in the filtered interferogram segments. When coupled with an aircraft-based application, the quantitative determination of target analytes by passive FTIR seems totally feasible.

Pattern Recognition Methodology Based on Filtered Interferogram Segments

Overview of Methodology. The digital filtering methodology described above effectively serves as a data preprocessing tool to reduce a collected interferogram to an interferogram segment that has been treated (i.e. windowed and filtered) to enhance the information regarding the presence of a targeted spectral band. The next processing step for a passive FTIR sensor is a decision-making step: Does the filtered segment indicate that the target analyte is present? If the analyte is present, an alarm is presumably triggered. If the analyte is not present, nothing is done.

Decision-making algorithms for numerical data are termed pattern recognition procedures. All such algorithms are based on two premises: (1) that one or more specific data categories or "classes" can be defined; and (2) that data vectors of representative members of each of these classes are available. In general, pattern recognition algorithms use this known, representative data to produce some type of mathematical construct (vector, model, etc.) that yields a numerical response when applied to an input data




Figure 29. Points 175-250 in six interferograms derived from the helicopter-based data run at 1000 ft.




Figure 30. The region from $920-970\text{ cm}^{-1}$ in the spectra derived from the interferograms in Figure 29. The SF_6 band increases in area from the top to the bottom of the plot.




Figure 31. Points 175-250 in the six interferograms after application of the SF_6 matrix filter.




Figure 32. Plot of the magnitude of the filtered interferogram segment vs. integrated area of the SF_6 spectral band. The line shown was derived from the data points by regression analysis.

vector or "pattern." The numerical response produced is interpreted to determine to which data class the input pattern belongs. For a two-class pattern recognition problem, this interpretation effectively yields a yes-no response. For a given target analyte in the passive FTIR application, two data classes exist, corresponding to target-containing interferograms and "clean" interferograms. For the work reported here, the input patterns are filtered interferogram segments.

All pattern recognition methods explore the multidimensional "pattern space" defined by the input patterns. The patterns selected as representative of the data classes (the "training set") are used to characterize the regions of this space "occupied" by the defined data classes. Two general approaches can be taken to this investigation.

In one method, the characteristic region of a given class can be modeled, producing a vector space model that spans the region of the class. Unknown patterns are fitted to (i.e. projected onto) the class models to determine if they lie in the same region of space. The SIMCA method of Wold¹⁶ is an example of this approach.

The second basic pattern recognition approach attempts to model the region between data classes, thereby establishing separating boundaries. These separating boundaries are termed discriminants, and the overall procedure is termed discriminant analysis¹⁷. The boundary is typically defined as an $(n-1)$ -dimensional hyperplane, where n is the dimensionality of the patterns. Mathematically, the hyperplane is defined by an n -dimensional vector, w , that is oriented normal to the surface of the hyperplane. The vector, w , is often termed a "weight vector." For a two-class problem, only one such vector is needed. Unknown patterns are assigned to a data class by determining on which side of the boundary they lie. This calculation is defined by

$$d = w^T x \quad (4)$$

where w^T is the transpose of w , x is the input pattern vector, and the scalar value, d , is termed the discriminant score. If w is scaled to unit length, the sign of d determines on which side of the boundary x lies. Iterative numerical optimization procedures are typically used to define w , given the input set of patterns deemed representative of the data classes.

For the passive FTIR problem, the variety of possible infrared backgrounds dictates that a large set of patterns be used if any serious attempt is to be made to define a pattern space that is truly representative. Thus, the selection of a specific pattern recognition approach must be based largely on the feasibility of the method for use with a large data set. Methods based on the construction of linear discriminants are more easily applied to a large data set than are the class modeling techniques. For this reason, these methods have been used in the feasibility studies reported here.

Construction of Linear Discriminants. Data from two ground-based and three helicopter-based experimental runs were assembled to define a representative set of interferograms for the pattern recognition feasibility study. A total of 449 SF₆-active and 551 inactive interferograms were selected, defining a total data set of 1000 interferograms. These interferograms were categorized as either SF₆-active or inactive based on inspections of the corresponding transformed spectra. Interferograms were sought with SF₆ levels ranging from very strong to the limit of detection. Since some interferograms with very weak SF₆ signals were sought, some discrepancies may exist in the actual classification of the 1000 interferograms as active or inactive.

For the feasibility study reported here, the pattern vectors used were defined in a straightforward manner. Points 175-250 in the interferograms

were filtered by use of the SF₆ matrix filter described previously. The filtered segments could not be used directly for pattern recognition, as the information pertaining to absorption and emission bands is 180 degrees out of phase (see discussion of Figures 9-11). This results in sign differences for the individual filtered points. For this reason, two data transforms for removing sign information were tested. As a first test, the values in the filtered segments were squared, employing a strategy similar to that used in the sum-of-squares vs. interferogram number plots presented previously. Second, the absolute value was taken across the filtered segment. In comparing the results of these transforms, the absolute value transform was found to yield the best results in the pattern recognition trials. It is hypothesized that this transform is superior, as it is a linear transform. As indicated by eq. 4, the discriminants developed here are also linear. The nonlinear effects of the squaring operation do not seem to be compatible with the use of linear discriminants.

The 1000 76-point transformed pattern vectors were submitted to a standard discriminant optimization procedure. This iterative procedure produced a weight vector that could place 981 of the 1000 interferograms in the correct class (active vs. inactive). The 19 misclassified patterns (9 active, 10 inactive) appear to define the limit of detection.

Testing of Computed Discriminant. The computed linear discriminant was tested by applying it to interferograms in several experimental runs. In each case, a plot was made of discriminant score vs. interferogram number. The discriminant score used for plotting was taken as the mean of the current computed score and the computed score of the previous interferogram. This effectively applies a two-point moving-average filter to the discriminant scores. This procedure decreases the incidence of false alarms by downweighting the effects of single interferograms on the alarm status.

The discriminant was first applied to two of the experimental runs from which interferograms were extracted for the discriminant development. Figure 33 is a plot of the computed two-point mean discriminant score vs. interferogram number for a ground-based run, while Figure 34 is a corresponding plot for a helicopter-based run. Those discriminant scores exceeding 0.0 would trigger an alarm for the presence of SF₆. Each of the interferograms for which an alarm is clearly indicated does appear to contain SF₆ information.

Figures 35 and 36, respectively, depict discriminant score plots for the helicopter-based and ground-based experimental runs used previously in the evaluation of filter performance. None of these interferograms were included in either the filter or discriminant development work. Again, the SF₆-active interferograms are clearly distinguished from the inactive interferograms. Although one false-alarm is seen in the helicopter-based run and two are seen in the ground-based run, this false-alarm rate represents only an 0.24% occurrence. Based on these results, standard pattern recognition procedures appear completely compatible with filtered interferogram segments.

Similar pattern recognition trials were then performed with pattern vectors formed from the reduced set of 51 points found during the matrix filter development work. The same set of 1000 interferograms was used, and the absolute value transform was applied as before. The best weight vector obtainable classified only 850 of the 1000 interferograms correctly, however. Subsequent application of this weight vector to the test data runs revealed much poorer discrimination than was observed with the 76-point discriminant.

Results obtained during the development of the matrix filter indicated that the 25 interferogram points deleted possessed low signal-to-noise for the sum of squares calculation. Such points may be useful for discrimination, however, if the active or inactive interferograms take on characteristic values for the points. This question was investigated by plotting the




Figure 33. Two-point mean discriminant score vs. interferogram number for a ground-based data run. Of these interferograms, 392 were used in the discriminant development.




Figure 34. Two-point mean discriminant score vs. interferogram number for a helicopter-based data run. Of these interferograms, 137 were used in the discriminant development.




Figure 35. Two-point mean discriminant score vs. interferogram number for a helicopter-based data run. None of these interferograms were used in the discriminant development.




Figure 36. Two-point mean discriminant score vs. interferogram number for a ground-based data run. None of these interferograms were used in the discriminant development.

percentage of the variance in the dependent variable explained in the matrix filter computation vs. the corresponding computed discriminant value for the point. This plot is depicted in Figure 37.

The discriminant values take on both positive and negative quantities, corresponding to additive and subtractive terms in the summation of the overall discriminant score (eq. 4). Nineteen of the 25 points corresponding to less than 90% variance explained have negative discriminant values. Of the six positive discriminant values, five are less than 0.1, giving them low overall weight in the determination of the discriminant score. The points with negative discriminant values contribute to lowering the discriminant score (i.e. leading to an inactive classification of the interferogram). In considering this point, note that the absolute value transform causes all values in the pattern vector to be positive.

These results confirm that for pattern recognition purposes, greater discrimination between active and inactive interferograms is found by including all of the points in the interferogram segment. In this example, the inactive interferograms appear to have large values at the points corresponding to low percentage variance explained. These large values would be expected to lower signal-to-noise in the sum-of-squares plot. For the pattern recognition procedure, however, they help to improve the discrimination between active and inactive interferograms. This illustrates the power of having both additive and subtractive information in a discriminant calculation. For this reason, discriminants typically outperform simpler procedures such as applying an alarm threshold to sum of squares plots of the type generated during the filter development work.




Figure 37. Plot of percentage variance explained in the matrix filter calculation vs. computed discriminant value for interferogram points 175-250.

Critical Evaluation of Time-Domain Methodology

The concepts described above for time-domain processing of passive FTIR data have been presented to several groups of knowledgeable FTIR spectroscopists. Two criticisms of the approach have been commonly raised and need to be addressed.

Questions of Spectral Resolution. It has been argued that the use of a short interferogram segment limits the analysis to only low-resolution spectral information. This argument is based on the standard FTIR principle that a higher resolution spectrum is obtained by collecting a longer interferogram. This argument is invalid on two grounds.

First, in conventional FTIR, a longer interferogram must be collected to obtain high resolution information due to limitations of the FFT algorithm. The FFT assumes the input waveform has the value zero outside the sampled region. Thus, the algorithm treats a short interferogram segment as a harmonic waveform that has been convolved with a boxcar function. The transform thus contains $\sin x/x$ components corresponding to the transform of the boxcar function. In practical terms, this results in a seriously broadened spectrum if the sampled interferogram segment is short. As an example, the top spectrum in Figure 38 is the result of applying the FFT to a 76-point interferogram segment. This interferogram consisted of a single cosine frequency at 941.1 cm^{-1} . The transform should yield a single spike at that frequency. The resultant spectrum is greatly broadened, however. The lower plot in the figure is the spectrum resulting from applying an alternative transform technique to the 76-point segment. The autoregressive transform algorithm of Marple¹⁸ attempts to model the frequencies in the input waveform. No assumption is made about the character of the waveform outside the sampled region. The resultant spectrum is very close to a single spike. A long interferogram is required only when the FFT is being used to extract high-resolution information. Arguments based on inherent characteristics of the FFT clearly do not apply when the FFT is not used.

Second, the discussion of Figures 9-11 clearly indicated that information pertaining to spectral features of different widths is present in different concentrations at different points in the interferogram. As shown in the figures, the information pertaining to the spectral band of interest is dominant when other features (e.g. the overall filter bandpass shape) have damped out. The selection of an appropriate interferogram segment is purely based on the spectral signature within the filter bandpass. Thus, in the time domain, questions of resolution are directed to location within the interferogram, rather than to the size of the interferogram segment. If a single spectral band is being extracted with a digital filter, the interferogram region of points 175-250 seems close to optimum for the typical widths of major bands in organic molecules.

Effects of Phase Errors. Phase errors result from deviations in the sampling points along the time axis of the interferogram. Most often, these are fractional point shifts that cause the interferogram to be slightly out of registration. In spectral-based processing, phase errors most often result in small spectral artifacts. A phase correction algorithm is typically applied in conjunction with the FFT to remove these artifacts. In time-domain processing, however, phase errors are potentially more damaging, as great dependence is placed on correct positioning within the interferogram. The use of single-point matrix filters places further demands on proper registration of the interferogram.

During the development work described in this report, no obvious deleterious effects due to interferogram phase errors have been observed. We are of the belief that most phase errors encountered are fractional point shifts that do not seem to have great effect. An experiment was conducted,




Figure 38. Transformed spectra corresponding to a 76-point segment taken from an interferogram composed of a single cosine frequency (941.1 cm^{-1}). Top: the spectrum resulting from application of the FFT. Bottom: the spectrum resulting from application of the autoregressive algorithm of Marple¹⁸.

however, to test the effects of catastrophic phase errors on the pattern recognition results. This experiment focused on the helicopter- and ground-based data sets whose discriminant score plots were presented in Figures 35 and 36, respectively. For each interferogram, a different point in the region of points 175-250 was selected randomly and deleted. The interferogram was then compressed to simulate a point being skipped while sampling the interferogram. The matrix filter and computed discriminant were then applied to the interferogram as if nothing untoward had happened. Figures 39 and 40 are the resulting discriminant score plots for the helicopter- and ground-based data sets, respectively.

The plots of discriminant scores are noisier than before, but the SF_6 -active interferograms are still clearly indicated. Several additional false alarms are now indicated in the ground-based plot, but overall, the serious phase errors introduced do not destroy the detection capability of the algorithm. When it is considered that the majority of phase errors encountered are less severe than those created here, it can be argued strongly that the current time-domain algorithm is sufficiently resistant to phase errors to be usable.

Source of False Alarms. In a final evaluation of the time-domain processing algorithm, interferograms generating false alarms were investigated to determine the cause of the false alarms. Those interferograms generating a discriminant score just above 0.0 tend to have a noisy spectral baseline across the filter bandpass. This causes the filtered interferogram segment to be similar to that obtained from a trace level of SF_6 . Those false alarms associated with discriminant scores significantly greater than 0.0 tend to result from some hardware malfunction during collection of the interferogram. The false alarm at interferogram 503 in Figure 35 is one such example. Figure 41 is the complete transformed spectrum of this interferogram. In processing this interferogram, the huge low-frequency noise spike seen in the spectrum




Figure 39. Two-point mean discriminant score plot for the helicopter-based data run after introduction of catastrophic phase errors.




Figure 40. Two-point mean discriminant score plot for the ground-based data run after introduction of catastrophic phase errors.




Figure 41. Transformed spectrum of interferogram 503 in Figure 35. The large low-frequency noise spike leads to a false alarm.

saturated the low-frequency cutoff filter used prior to the matrix filter. A significant amount of low frequency noise was left in the interferogram, ultimately leading to the false alarm. This example illustrates a problem with time-domain processing. If the digital filters in use break down, false alarms may result. This interferogram would likely not give rise to a false alarm in a spectral-based processing scheme. Some means for detecting these catastrophic phenomena needs to be included in the time-domain processing to prevent false alarms of this type from occurring.

CONCLUSIONS

The results presented in this report confirm the practicality of an interferogram-based processing procedure for passive FTIR data. The development of matrix filters represents a fundamentally new approach to signal processing of time-domain data. The high accuracy of the filters allows the center and width of the filter bandpass to be determined precisely. In addition, easing the demands on individual filters reduces the number of coefficients within those filters, resulting in large savings in computation time. In this regard, it should be stressed that the savings in computation time reported in Table III assume that a spectral-based analysis requires only an FFT. In reality, spectral phase correction would be used, along with some type of spectral subtraction and/or digital filtering. Therefore, the quoted savings in computation time are conservative.

The significance of this work can be seen in two areas. First, the ability to use only a small set of discrete interferogram points makes possible a new generation of passive FTIR sensors in which only a short interferogram segment is collected. Such an instrument could be built around a simpler interferometer than is currently used, thereby increasing the reliability of the sensor and decreasing its cost. The reduction in

computational requirements simplifies the required computer hardware in an analogous manner.

The second point of significance to the work is the elimination of the requirement for subtraction of a background spectrum. The two target analytes used, SF₆ and DMMP, were both detected under rapidly changing infrared backgrounds. The filter-based processing strategy should therefore feature increased resistance to background changes over a spectral subtraction approach.

In summary, the research described here provides new capabilities for the analysis of data from passive FTIR sensors. When coupled with recent advances in FTIR instrumentation, it should now be completely feasible to configure a small, economical toxic gas alarm system for both stationary and mobile monitoring applications.

RECOMMENDATIONS

In effect, this report describes a feasibility study directed at evaluating the potential of an interferogram-based processing scheme for passive FTIR data. Without question, the work indicates that such a scheme is feasible. For the algorithm to be made ready for actual implementation, however, three areas of further investigation are most important.

First, the resistance of the methodology to chemical interferents must be addressed. In this regard, we refer to interferents whose spectral features lie within the filter bandpass of a given target analyte. Testing of the methodology with a variety of analyte/interferent combinations is essential.

Second, a formalized optimization study must be performed to establish which interferogram points should be used in the time-domain processing. Work performed to date indicates that the optimum interferogram segment depends on the center and width of the filter bandpass. Therefore, for each analyte to be detected, these three variables must be optimized. The development of formalized procedures for performing this optimization is critical to the optimization of the overall time-domain processing scheme.

Lastly, for the discriminant development work, the effects of training set size must be addressed. Practical limits need to be established for the necessary size of the training set. In addition, an algorithm for the selection of training set members must be developed. These considerations are very important if the discriminant values are to be firmly established.

In conclusion, studies such as those outlined above will serve to refine the methodology. We are of firm belief that the time-domain processing concepts are sound. Work must now proceed to optimize the specific procedures used.

Blank

REFERENCES

1. Hanst, P. L. in "Advances in Environmental Science and Technology", Pitts, J. N., Metcalf, R. L., Eds.; Wiley: New York, 1971.
2. Chan, S. H.; Lin, C. C.; Low, M. J. D. Environ. Sci. Technol. 1973, 7, 424-428.
3. Walter, H.; Flanigan, D. Appl. Opt. 1975, 14, 1423-1428.
4. Low, M. J. D.; Clancy, F. K. Environ. Sci. Technol. 1967, 1, 73-74.
5. Herget, W. F.; Brasher, J. D.; Opt. Eng. 1980, 19, 508-514.
6. Mertz, L. U. S. Dept. of Commerce Rept. No. AD602936, p. 120.
7. Block, L. C.; Zachor, A. S.; Appl. Opt. 1964, 3, 209.
8. Hanel, R. A.; Schachman, B.; Clark, F. D.; Prokesh, C. H.; Taylor, J. B.; Watson, W. M.; Chaney, L. Appl. Opt. 1970, 9, 1767.
9. Herget, W. F. Abstracts of Papers, SPIE 1981 International Conference on Fourier Transform Infrared Spectroscopy, Columbia, SC; pp 449-456.
10. Kauffman, J. W. Abstracts of Papers, SPIE 1981 International Conference on Fourier Transform Infrared Spectroscopy, Columbia, SC; pp 426-435.
11. Walker, R. A.; Rex, J. D. Proc. Soc. Photo-Opt. Instrum. Eng. 1979, 191, 88.
12. Childers, D.; Durling, A. Digital Filtering and Signal Processing; West Publishing: St. Paul, MN, 1975.
13. McClellan, J. H.; Parks, T. W. IEEE Trans. Circuit Theor. 1973, CT-20, 697-701.
14. McClellan, J. H.; Parks, T. W.; Rabiner, L. R. IEEE Trans. Audio Electroacoust. 1973, AU-21, 506-525.
15. Draper, N. R.; Smith, H. Applied Regression Analysis, 2nd ed.; Wiley-Interscience: New York, 1981.
16. Wold, S.; Sjostrom, M. in Chemometrics, Theory and Practice, B. R. Kowalski, Ed.; American Chemical Society: Washington, DC, 1977.
17. Nilsson, N. J. Learning Machines; McGraw-Hill: New York, 1965.
18. Marple, L. IEEE Trans. Acoust., Speech, Signal Processing 1980, ASSP-28, 441-454.